MobiCom Poster: Practical Experience with Wireless Networks Integration using Mobile IPv6

Rajiv Chakravorty[a] **Pablo Vidales**[b] **Kavitha Subramanian**[a] **Ian Pratt**[a] **Jon Crowcroft**[a]
{*Rajiv.Chakravorty, Pablo.Vidales, Kavitha.Subramanian, Ian.Pratt, Jon.Crowcroft*}*@cl.cam.ac.uk*
[a] University of Cambridge Computer Laboratory
[b] Laboratory for Communication Engineering
William Gates Building, JJ Thompson Avenue, Cambridge CB3 0FD, U.K.
Project Web: http://www.cl.cam.ac.uk/coms/

I. Introduction

Efficient integration of heterogeneous wireless access networks is a key step in building an all-IP ubiquitous wireless access infrastructure. Mobile IPv6 can play a key role in integrating different link-layer technologies, with the promise of enabling *transparent* mobility through use of a unified network layer. In Cambridge Open Mobile System Project (`http://www.cl.cam.ac.uk/coms/`), we are investigating:

- partitioning the latency components in a Mobile IPv6-based handover,
- identifying TCP and HTTP performance problems during vertical handovers using Mobile IPv6,
- exploring the extent to which Mobile IPv6 protocol can hide differences of the disparate underlying link-layer technologies, and,
- demonstrating the efficacy of *multi-layer optimization* schemes that can improve performance during vertical handovers.

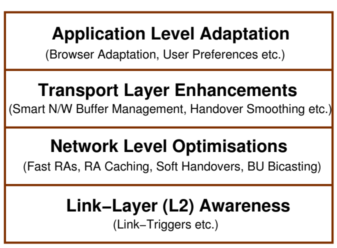

Figure 1: Handover Optimizations

II. Test Environment and Tools

Our experimental setup consists of a Mobile IPv6-based loosely-coupled LAN-WLAN-GPRS testbed as shown in figure 2. In this testbed, the cellular GPRS network infrastructure currently in use is the Vodafone UK's production GPRS network. The WLAN access points (APs) are IEEE 802.11b APs located at different locations of the building.

The GPRS infrastructure comprises base stations (BSs) that are linked to the SGSN (Serving GPRS Support Node) which is then connected to a GGSN (Gateway GPRS Support node). A well provisioned virtual private network (VPN) connects the Lab network to that of the Vodafone's backbone via an IPSec tunnel over the public Internet. A separate "operator-type" RADIUS server is provisioned to authenticate GPRS mobile users/terminals and also assign IP addresses.

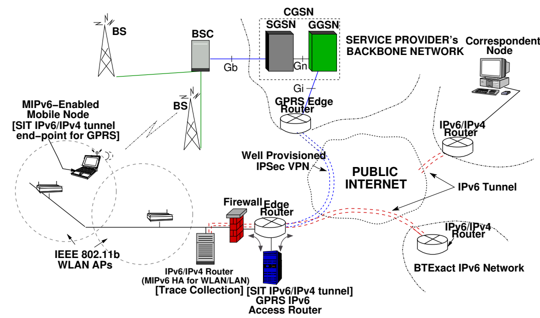

Figure 2: GPRS–WLAN–LAN Testbed.

We brokered a semi-permanent IPv6 subnet from BTExact's IPv6 Network, which connects us to the 6BONE. Using the address space, we are able to allocate static IPv6 addresses to all our IPv6 enabled mobile nodes. A router in the lab acts an IPv6/IPv4 tunnel end-point to the BTExact's IPv6 network (shown in figure 2). This router is also an IPv6 access router (Home Agent) for the lab's fixed-internal IPv6-enabled network and also for internal WLANs. Routing in the Lab has been configured such that all GPRS/WLAN user traffic going to and from mobile clients are allowed to pass through the internal router, enabling us to perform traffic monitoring.

All the characterization tests for GPRS-WLAN-LAN vertical handovers were analysed using a version of `tcptrace` program updated (`tcptrace+`) to trace TCP connections for Mobile IPv6 handovers.

III. Experimental Evaluation

We investigated the extent to which Mobile IPv6 could be used to successfully migrate TCP connections during inter-network handovers. From the handover characterization process, we split the Mobile IPv6 handover latency into three main components – detection time (t_d)(includes movement and any duplicate address detection time), configuration time (t_c), and registration time (t_r), each of which contribute to the overall handover latency [1].

We studied the effects of mobility on vertical handovers, and have highlighted the challenges with IP mobility. Using the testbed, we have evaluated the impact layer-3 *hard* handovers have on transport protocols such as TCP. Here, we summarize our practical experiences, a thorough description is available in the form of a separate technical report [1].

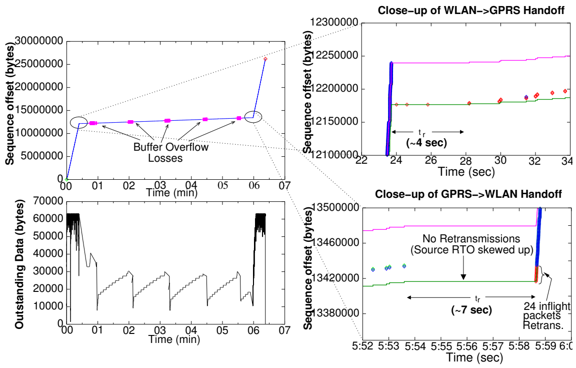

Figure 3: tcptrace+ of WLAN↔GPRS handover.

We conducted vertical handovers between different networks – GPRS↔WLAN and GPRS↔LAN – using a multimode mobile device located in a WLAN hot-spot (and LAN) and also under GPRS coverage.

Figure 3 shows a typical TCP connection behaviour during a vertical handover between GPRS↔WLAN. We find that the time to handover from GPRS→WLAN and WLAN→GPRS to be high (7s and 4s respectively). Consequently, TCP connections experience multiple time-outs before the handover is complete. A number of factors contribute to such high handover latency, but mainly the disparity in the link-layer characteristics between the two networks (see [1] for more information). [1] also gives mean latency of each handover component in t_d, t_c, and t_r.

IV. Summary of Optimizations

We have evaluated schemes for handover optimization specificallly for improving TCP performance. Here, we briefy discuss our practical experiences with these optimizations.

`Hard handover optimizations:` We applied the following optimization schemes to improve hard handover performance:

- `Fast Router Advertisements (Fast RAs).` We used Fast RAs to improve handover performance, and have shown that reducing RA interval also reduces detection time during handovers. However, we also observed that reducing RA interval to very low values (40ms-70ms) as that typically specified by the latest Mobile IPv6 draft leads to substantial overhead in GPRS (upto 25-50% of total bandwidth), and does not always guarantee significant improvement in mean detection times.

- `Client-based RA Caching.` RA caching is a technique to eliminate detection time during vertical handovers. In this scheme, RAs are cached a *priori* by the mobile client, so that when the decision to handover is taken, the detection time for RA lookup during handover execution is eliminated, thereby improving handover performance. We implemented the RA cache as a Linux 2.4 module, and have shown that the benefits of eliminating detection time ($t_d \sim 0$) during vertical handovers is significant.

- `Smart Buffer Management using TCP Proxy.` We have shown that excess buffering due to long (or a large number of) TCP connections in the GPRS GGSN node can inflate the source TCP's RTT, and hence, skew the retransmission timer (RTO). This can prevent the source re-transmitting during handover from GPRS→WLAN. Current GPRS networks offer excess buffering, and this can be prevented using smart buffer management. We used a TCP enhancing proxy that prevents the source from excessively buffering data, and have shown substantial improvement in registration times during handovers.

- `Client-Assisted Binding Update Simulcast.` We observed that when handing off from WLAN→GPRS, the registration process of binding updates (BU) to the correspondent node would entail the high RTT of the GPRS link. BU simulcasting is a scheme to bicast BUs not only from GPRS, but also from other fast networks (e.g., WLAN/LAN) just before a client decides to handover. We implemented BU simulcast as an extension to the MIPL Mobile IPv6 source code and have shown that it can further reduce registration times during *upward* vertical (WLAN→GPRS) handovers.

`Layer-3 based Soft handovers:` Unlike hard handovers, a better technique to handover across coverage offered by networks in a wireless overlay can exploit the inherent macro-diversity available in order to make handovers *soft* to help improve performance. We implemented the *soft* handover scheme as a Linux module (along with RA caching) and figure 4 shows one such trace for TCP using soft handovers with RA caching enabled. We have found TCP performance to dramatically improve with soft handovers (and RA caching enabled) with only 0.7s required for the handover than the one we have shown earlier (with hard handover in figure 3). Traditionally, soft handovers have been successfully exploited for link-layer handovers in cellular networks.

These experimentations reveal the efficacy of the soft handover approach, which when applied with other (hard) handover optimization schemes as discussed earlier, improves TCP performance during vertical handovers dramatically.

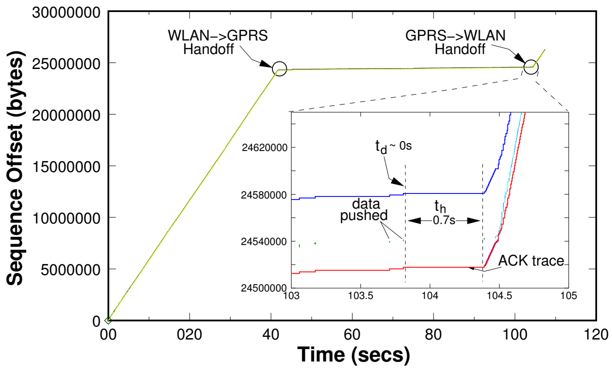

Figure 4: Exploiting Overlay diversity in (Soft) handovers.

V. Open Issues for Research

There are many issues resulting from this research. Specifically, we ask the following questions:

1. How do the benefits quantify at different layers in the handover optimization schemes?
2. Do schemes that benefit TCP during handover imply commensurate benefits to HTTP (web) performance? How does that translate quantitatively for HTTP?
3. Can soft handover approach applied at Layer 3 be used for high mobility environments? What quantitative benefits can we achieve using a similar approach even for real-time flows?
4. How can we have smooth handovers while handing off from faster/fatter WLAN to slower/thinner GPRS?
5. In soft handovers, can we prevent the source TCP from entering fast-restransmit mode (due to dupACKs generated by the mobile client) during handover from slow (GPRS) to fast networks (WLAN)?

COMS website `http://www.cl.cam.ac.uk/coms/` provides further information about our ongoing research and papers.

References

[1] R. Chakravorty, P. Vidales, L. Patanapongpibul, K. Subramanian, I. Pratt and J. Crowcroft. "On Inter-network Handover Performance using Mobile IPv6". University of Cambridge Computer Laboratory– Technical Report, June 2003. `http://www.cl.cam.ac.uk/coms/publications.htm`